From Health Search to Health Care: Explorations of Intention and Utilization via Query Logs and User Surveys

Ryen W. White, Ph.D. and Eric Horvitz, M.D., Ph.D.
Microsoft Research
One Microsoft Way
Redmond, WA 98052 USA
{ryenw, horvitz}@microsoft.com

Keywords: Healthcare utilization, healthcare seeking, search log analysis, user surveys.

Word count: 4864

ABSTRACT

Objective

We seek to better understand the relationship between online health-seeking behaviors and in-world healthcare utilization (HU). We do so via studies of online search and access activities before and after queries that pursue medical professionals and facilities.

Materials and Methods

We analyze data collected from logs of online searches gathered from consenting users of a browser toolbar from Microsoft (*N* = 9,740). We employ a complementary survey (*N* = 489) to seek a deeper understanding of information-gathering, reflection, and action on the pursuit of professional healthcare.

Results

We provide insights about HU through the survey, breaking out its findings by different respondent marginalizations as appropriate. Observations made from search logs may be explained by trends observed in our survey responses, even though the user populations differ.

Discussion

The results provide insights about how users decide if and when to utilize healthcare resources, and how online health information seeking transitions to in-world HU. The findings from both the survey and the logs reveal behavioral patterns and suggest a strong relationship between search behavior and healthcare utilization. Although the diversity of our survey respondents is limited and we cannot be certain that users visited medical facilities, we demonstrate that it may be possible to infer HU from long-term search behavior by the apparent influence that health concerns and professional advice have on search activity.

Conclusion

Our findings highlight different phases of online activities around queries pursuing professional healthcare facilities and services. We also show that it may be possible to infer HU from logs without tracking people's physical location, based on the effect of HU on pre- and post-HU search behavior. This allows search providers and others to develop more robust models of interests and preferences by modeling *utilization* rather than simply the *intention to utilize* that is expressed in search queries.

BODY

Background and Significance

A recent survey by the Pew Internet and American Life Project found that 59% of U.S. adults have looked online for health information in the past year, primarily through search engines such as Google, Bing, and Yahoo! (Fox and Duggan, 2013). In analysis of survey responses, Sillence and colleagues (2006) found that over 73% respondents used the Web for health advice, support, or preparation for an appointment. In many situations the online pursuit of health information may be a critical phase in the transition from medical concerns to utilizing professional medical care. Research has shown that Web content can drive people to visit their general practitioner or medical specialist (Baker et al., 2003; White and Horvitz, 2009a). In the Pew survey cited above, 35% of respondents diagnosed a medical condition online, and of those people, just over half of them followed up their online activity with a visit to a medical professional.

Research on consumer health information on the Web (Cline and Haynes, 2001; Eysenbach and Kohler, 2002) shows that medical professionals should be concerned about both the quality of online health information, and people's apparent reluctance to consider source quality in making medical decisions. Other studies have examined the relationship between the review of online health information and HU. Ayers and Kronenfeld (2007) explored the relationship between chronic medical conditions and frequency of Web use, as well as changes in health behavior associated with Web usage. They found that the more frequently someone uses the Web as a source of health information, the more likely they will alter their healthcare behavior. Eastin and Guinsler (2006) showed that an individual's level of health anxiety moderates the relationship between online health information seeking and HU decisions. Others have investigated signs of the pursuit of information on healthcare facilities and providers from search logs (White and Horvitz, 2010) and used privacy-sensitive mechanisms to study in-world medical search behavior from queries tagged with global positioning system (GPS) coordinates (White and Horvitz, 2013). More generally, search log data has been used to study how people perform search (White and Drucker, 2007), to predict their next online actions (Lau and Horvitz, 1999; Downey et al., 2007), to predict their future interests (Dupret and Piwoworksi, 2008), to improve search engines (Joachims, 2002; Tan et al., 2006), and to understand in-world activities from long-term logs (Richardson, 2009).

These previous studies have not examined the processes by which people *transition* from Web search to seeking professional medical care, including factors that affect their decisions about which facility to select and changes in their online search behavior that can help predict HU events. The rationales behind HU decisions, and the relationship between HU and Web search (before and after the consultation) are poorly understood. In this study we show a relationship between the timing of when HU events are likely to occur (from a survey) and patterns of search behavior (mined from search-engine query logs). Estimating HU directly from search behavior can help search providers develop richer user interest models for applications such as search-result personalization, depending on the phases of health concern and information seeking (e.g., support diagnosis before the (predicted) visit, assist with the transition from Web search to HU, and help with understanding conditions and treatments post visit). Inferring HU indirectly from its impact on search behavior has the welcome side effect that users' physical locations need not be tracked (e.g., with GPS sensors (White and Horvitz 2013)) to establish that professional care has been sought.

Objective

In this study, we aim to better understand the rationale behind people's HU decisions, including when they decide to visit, why they visit, and which facilities they select. We also seek to understand the relationship between such HU activities and Web search, specifically: (1) the role the Web plays in helping people decide when and where to seek professional medical attention, and (2) the impact that visits might have on Web search behavior both before and after a visit, highlighting the potential of search behavior as an indirect sensor of HU. Better understanding the choices and search behaviors surrounding HU can be used for a range of personalization, recommendation, and forecasting tasks.

Materials and Methods

We employ two complementary methodologies: (1) analyses of search-engine log data to study search behavior, and (2) the use of a survey to understand people's perceptions and motivations when making HU decisions, and help explain trends noted in the logs.

Log Analysis

We studied the search logs of a large cohort of randomly selected medical Web searchers. We analyzed four months of anonymized logs of URLs visited by users who consented to provide their log data through a browser toolbar widely distributed by Microsoft, gathered from January 2011 to April 2011. Log entries included a unique user identifier, a timestamp for each Web page view, and the URL of the page visited. We excluded intranet and secure (https) URL visits at the source. To reduce variability associated with cultural and linguistic variations in search behavior, we only included log entries generated by users in the English-speaking United States locale. From these logs, we extracted queries to Google, Yahoo!, and Bing.

The anonymous user identifier allows us to study users' search behavior across multiple sessions and examine changes in their medical search behavior over time. We are particularly interested in identifying evidence of HU in the logs. Since we do not have the ability to confirm in-world actions of people performing online searches, we utilize a surrogate representation that provides evidence of HU from log data. Specifically, we identify queries that show HU *intent* (HUI) (White and Horvitz, 2010). Such queries provide evidence of the pursuit of in-world professional care, e.g., [dermatologist 98033] and [neurologist decatur, il]. Identification of these queries can be performed automatically by spotting terms or phrases describing physical locations (e.g., "hospital", "clinic") or medical professionals (e.g., "physician", "cardiologist").

We automatically annotated search queries based on whether they contained *symptoms* (e.g., headache), *benign explanations* (e.g., caffeine withdrawal), *serious illnesses* (e.g., brain tumor), or combinations thereof. We leveraged three manually curated lists:

- *Symptoms*: Symptom list from the Merck medical dictionary, used in previous analysis of online search behavior (Cartright et al., 2011) to identify health-related sessions.
- *Benign explanations*: Commonly-occurring conditions, as defined in White and Horvitz (2009b) and used in that study for a log-based analysis of online search behavior. The word list is based on the International Classification of Diseases 10th Edition (ICD-10) from the World Health Organization.

- *Serious illnesses*: Serious conditions, again from ICD-10, and used by White and Horvitz (2009b).

All query labeling was performed automatically using keyword spotting. We also used synonyms of symptoms and conditions to improve query coverage (e.g., including "tiredness" in addition to "fatigue"). Synonyms for each symptom or condition were identified via a two-step walk on the search engine esult click graph (Beeferman and Berger, 2000). The authors reviewed the automatically generated lists of synonyms to remove erroneous list entries (e.g., all astrology-related synonyms were removed for cancer). We avoided substring matches to ensure high precision. The resultant distribution of query types comprised symptoms (38%), serious illnesses (43%), and benign explanations (16%), with the remainder comprising different combinations of these three classes.

Users with health concerns may be at different stages in their seeking of health information (e.g., some performing early phases of self-diagnosis, whereas others actively seek care for professionally diagnosed illnesses). To address this, we identified the first evidence of searching with HU intent in the logs (i.e., the first HUI query) and aligned users by this landmark search.

Survey

To study people's perceptions of the process by which they reflect about and then pursue healthcare resources, we created a survey designed to elicit information on how people search for medical information on the Web and how that relates to visits to medical facilities. The survey was anonymous, contained a mix of approximately 50 open and closed questions covering a broad range of issues around HU, including how people located medical facilities, the scheduling of consultation or treatment at the locations, and Web search behavior before and after the visit. To our knowledge, the survey respondents did not include users who had contributed data in the online log study. We pre-tested the instrument with five volunteers and iterated on wording to help ensure question clarity. The survey was distributed and analyzed electronically. Responses were stored on a secure server with no association between survey responses and respondent identities.

Results

Log Analysis

Using the log data described earlier, we constructed profiles for a randomly selected subset of users who visited at least one URL labeled with the "Health" category of the Open Directory Project (ODP, dmoz.org). We aligned the search streams by the first observed HUI for each user and assessed changes in people's medical search behavior before and after this reference HUI. To ensure adequate coverage (as users may be issuing HUIs at different stages), we studied behavior for 50 days before the first HUI and 80 days after. Beyond those thresholds the data originated from fewer than 100 searchers, raising concerns about data sparsity. To improve the likelihood that the queries on symptoms and conditions are related to the HUI, we focus on queries referring only to the symptoms and conditions that appear in the *same session* as the HUI query. This filtering yielded 9,740 long-term profiles.

Figure 1 reports changes in the frequencies that we observe for symptoms, benign explanations, and serious illnesses identified using the methods described earlier. The first HUI is marked in the figure. The figure is zoomed into the period 20 days before through 50 days after the first HUI since we wanted to

highlight trends in search behavior around that time. A figure summarizing trends across the larger studied time period (–50 days to +80 days) is shown in the Appendix.

We computed the change in the number of people searching for queries with different types of medical concerns versus background (expected) activity across all searchers using the approach of Richardson (2009). Let $n(\delta, q)$ denote the number of users who queried for query *class* q, δ days before or after the *reference* (their first) HUI query. We focus on four classes of query: (1) queries for medical symptoms, (2) queries for benign medical explanations, (3) queries for serious medical illnesses, and (4) queries for HUIs (search for medical resources, appearing *after* the reference HUI). Also let $n(\delta)$ denote the number of users who queried for anything δ days from the first time they queried for an HUI (i.e., not conditioned on q and including non-medical topics). Given these user counts on each day before and after the reference HUI query, we consider the observed and expected numbers of unique people on a daily basis who issue different classes of query. $Observed = n(\delta, q)/\sum_{\delta} n(\delta, q)$, the number of people querying for q on day δ, normalized by sum of the number of unique people performing this query over the period of the study, where people are counted as active on a daily basis. $Expected = n(\delta)/\sum_{\delta} n(\delta)$, is the background activity level on each referenced day, reflecting all queries (medical and non-medical) over time. To quantify the extent of the increase for each query class q over the background, we compute a score for each day as $L(\delta) = Observed/Expected$. We shall refer to this score as the *lift,* as such increases over background expected rates are typically referred to in studies of online search behavior (Richardson, 2009). $L(\delta)$ is shown in Figure 1 for symptoms, benign explanations, and serious illnesses. The figure also shows $L(\delta)$ for follow-on HUI queries, which is computed in a similar way as for symptoms and conditions, but focuses on queries for medical resources such as physicians, hospitals, or medical specialists. We show these as a bar chart to clearly distinguish queries for HUIs from symptom/condition queries since they reflect a different type of medical search intention. To improve graph readability, we smoothed $L(\delta)$ using Gaussian-weighted averaging with a standard deviation of seven days.

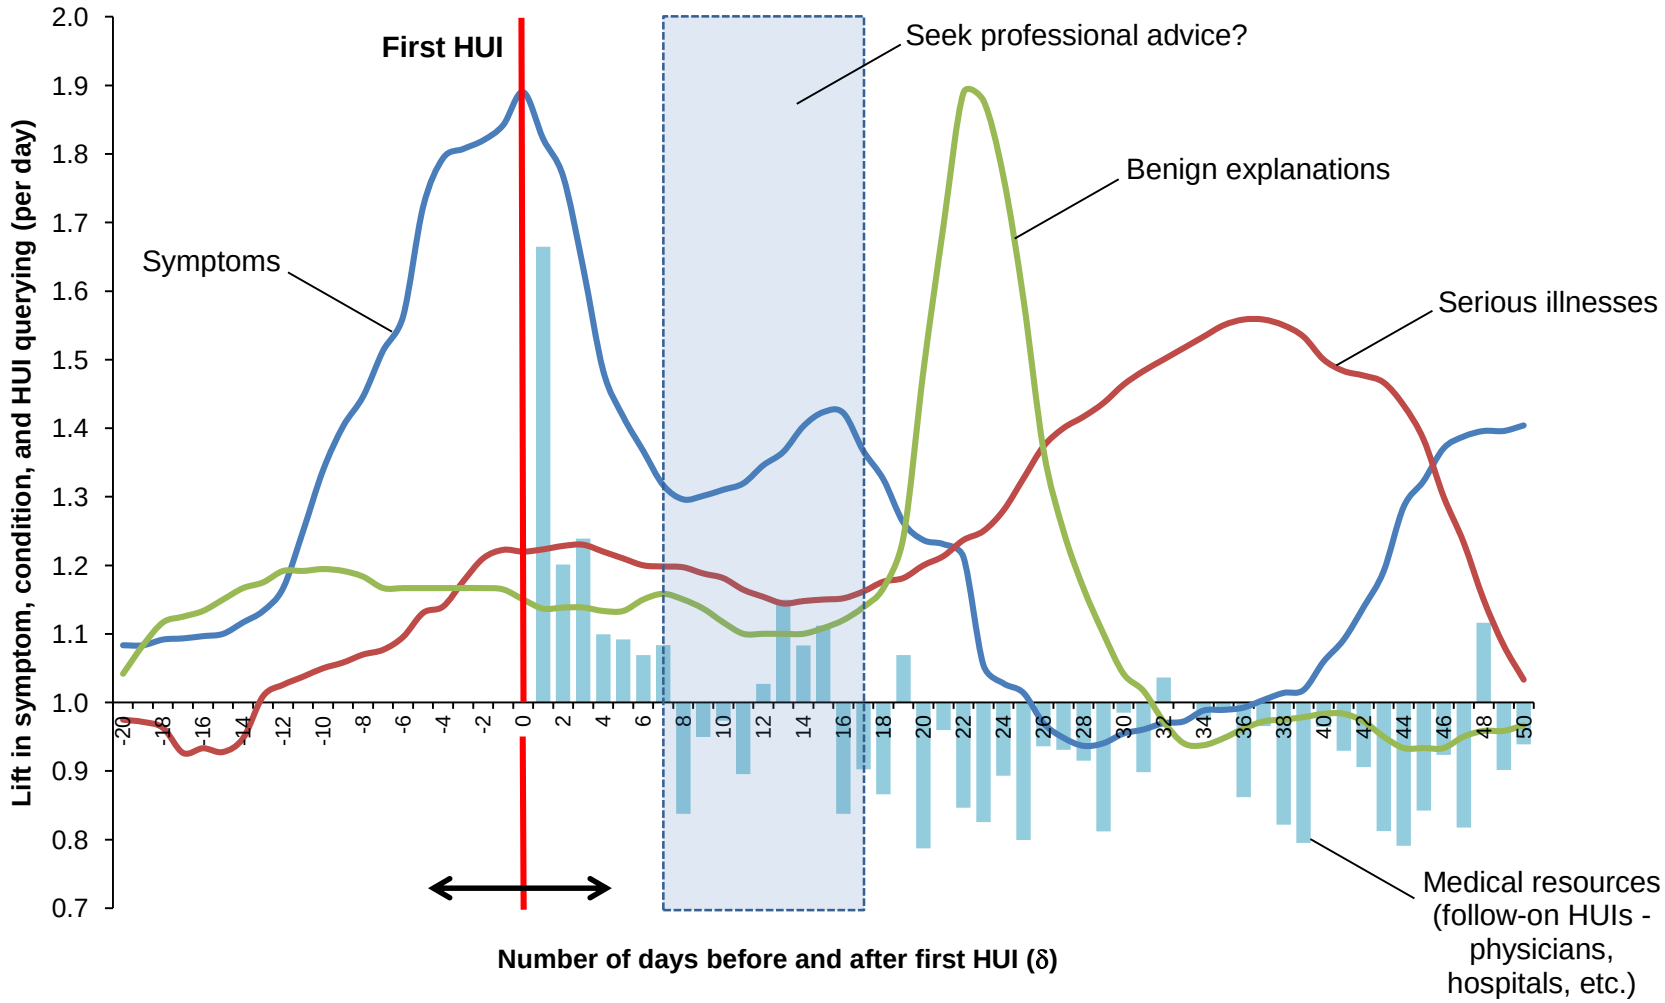

Figure 1: Increases over background, expected rates of each of the classes of query (symptoms, conditions, HUIs) over the background search activity (expected) for each day, with the first HUI as reference point. To reduce noise we focus on symptoms/conditions appearing in the same session as the reference HUI query. Bars depict the lift in interest in medical resources (healthcare professionals, medical facilities, etc.) over background search activity. The full (zoomed out) version of the figure is shown in the Appendix.

We make the following observations from Figure 1:

- A significant increase in symptom-related searches is observed adjacent to the occurrence of the first query with HUI, perhaps reflecting a heightened state of concern or uncertainty. This drops following the HUI, perhaps as people seek professional advice. We show in the survey presented later that this pursuit can happen almost immediately.
- A spike is observed in searches on common, typically-benign explanations for symptoms three weeks after the first HUI search, perhaps associated with patient calming from reassurance received during the meeting with the medical professional. The survey responses presented later in the article support this inference in two ways, *time* and *type of change*:
    - **Time:** Reponses on the time between the first HUI query and the visit (Table 1, Question 4) show that most respondents (94%) sought professional advice within a few weeks.
    - **Type of change:** Post-visit statistics reported in Table 3 show that many users reported an increase in benign explanation searching following their appointment.

  Following the spike in Figure 1, the amount of benign-explanation searching returns to a level similar to that before the concern emerged.
- Queries on serious illnesses increase a few weeks after the first HUI, perhaps also reflecting a more concerning diagnosis received during a visit.
- The number of symptom-related searches correlates well with the number of HUI searches (Pearson's correlation coefficient, $r = 0.64$, $t(78) = 14.43$, $p < 0.001$). Symptom searching is linked to uncertainty, and it seems reasonable that people would seek to resolve this uncertainty via medical consultation when it arises.
- HUI searches decrease after the first HUI, but reemerge about a week later (perhaps associated with confirming the visit, seeking directions to the facility, etc.), and then decrease to levels suggesting that users are then less likely than expected to issue HUI queries thereafter.

Overall, we see several noteworthy trends in search behavior surrounding the appearance of a HUI. In the next section we examine responses to our complementary survey where we asked people to recall their last visit to a medical facility and their behavior before and after. This helps us to help better understand healthcare utilization and the role that Web search plays in HU decisions.

Survey Analysis

In pursuit of additional insights to help explain and validate findings in the search logs, we administered a patient survey completed by those who reported seeking professional medical care. During Fall 2011, we sent an email invitation to 2,000 randomly selected employees within Microsoft's extended campus in Redmond, Washington. In total, 489 volunteers, who had indicated in a pre-screening that they could remember their last visit to a medical facility, completed the survey. No compensation was provided to respondents. For the survey, we consider all responses plus: (1) respondents' medical domain knowledge (high or low), and (2) whether it was the first time they visited the facility in question. Most respondents (72%) assessed that they had average levels of medical domain knowledge, with others reporting it to be high (13%) or low (15%). One in four respondents reported that the visit referred to in their pre-screening referred to the first time that they had visited the medical facility. To determine the statistical significance of observed differences we use Chi-squared tests and Z-tests on proportions between (i) *all visits* vs. *first*

visit and (ii) *high domain knowledge* vs. *low domain knowledge*, and test for significance at $p < .05$ and $p < .01$, using Bonferroni corrections as appropriate to counteract possible Type I errors.

Nature of the Visit to the Medical Facility

We asked respondents to provide general background information about the type of facility that they visited, why they visited, who the visit was for, and the distance that they traveled. Respondents reported mainly visiting doctors' offices (65%), and that their motivations were primarily medical (62%) and mostly related to symptoms (45%). Most participants had visited the facility on at least one previous occasion (75%), most of the visits were for themselves (81%), a loved one (18%), or a friend/colleague (1%). Over 95% of respondents reported that the facility that they visited was within 20 miles of their starting location (home, work, etc.). We will show later that proximity is an important determinant of the facility visited.

Locating Medical Facilities

We now focus on how people identified medical facilities and the role of Web searching in those decisions. Table 1 summarizes responses to questions about whether people searched for the facility prior to the visit and why they searched. The findings in the table (Q1) show that almost 40% of participants searched online for information about the facility prior to the visit. The main reasons given for searching were finding directions or information about physicians and the services offered (Q2). Respondents reported that the visits to the medical facility occurred typically within a week of the first search for that facility (78%) and nearly always within two weeks (94%). This region is marked in Figure 1 with "Seek professional care?". Interestingly, most participants still conducted follow-up searches after identifying the facility of interest (Q5: 61% Yes vs. 39% No; one-proportion Z-test: $Z = 4.93$, $p < 0.001$). The primary focus was on comparing different facilities (51%) and comparing physicians (42%).

Considering the conditionalizations, first-time visitors were more likely to search for information before the visit (Q1: 66% First-time visits vs. 38% All visits), and were also more likely to search for the facility to find the medical services that were offered (Q2: 35% First-time visits vs. 24% All visits). We also observed differences attributable to knowledge level. Searchers with low domain knowledge were more likely to search before the visit (Q1) and more likely to search for a *type* of facility rather than a particular facility (e.g., [clinics near casper montana]) (Q4). Those with high domain knowledge were likely to visit either immediately (35%) or after 1-2 weeks (40%). In contrast, the visits of those with less domain knowledge primarily visited within the week following the HUI query (85%), with 60% visiting within one day. Much fewer (only 15%) of these low-expertise users visited after two weeks. One explanation is that those with higher domain knowledge can make better decisions about the urgency of pursuing treatment associated with observed symptoms. They may be more capable of distinguishing between ailments that need urgent care, and those for which professional care can be postponed.

Table 1: Breakdown of reasons provided by participants in seeking medical facilities. Statistical significance between the groups of interest is denoted by ○ $p < .05$; ● $p < .01$.

Question and Response Options	All Visits (*N*=489)	Conditional responses		
		First Visit (*N*=122)	Domain Knowledge	
			High (*N*=63)	Low (*N*=75)
Q1. Did you search online for information about the medical facility that you visited prior to the visit?				
Yes	37.8%	65.6%	36.5% ○	45.3%
No	62.2% ●	34.4%	63.5%	54.7%
Q2. **If Yes**, why did you need information? (Select one or more reasons)	***N*=185**	***N*=80**	***N*=23**	***N*=34**
Find a new doctor or specialist	37.3%	45.0%	26.1% ○	32.4%
Find the medical services offered	24.3% ●	35.0%	13.0% ●	32.4%
Visitor information (e.g., visiting hours)	21.6%	22.5%	17.4% ●	29.4%
Directions	45.4%	48.8%	47.8%	47.1%
Phone number	15.7%	10.0%	8.7% ●	20.6%
Employment (e.g., job vacancies)	0.5%	0.8%	0.0%	0.0%
Educational (e.g., classes offered by the facility)	1.6%	1.3%	4.3%	0.0%
Other (please specify)	10.3%[1]	11.2%	13.0%	8.8%
Q3. What was the approximate time from the first time you searched for information on the medical facility until your visit?				
Immediately	25.8%	25.8%	35.0% ○	25.9%
1 day	25.8%	27.3%	15.0% ●	33.3%
1 week	26.5%	27.3%	20.0% ○	25.9%
2 weeks	15.9%	15.2%	20.0% ●	3.7%
Longer	6.0%	4.5%	10.0%	11.1%
Q4. Did you search for information specifically for the facility that you visited (versus, say, more generally for facilities of this *type*)?				
Yes	71.9%	72.5%	78.3%	55.9% ●
No	28.1%	27.5%	21.7%	44.1%
Q5. Did you visit the medical facility or professional that you searched for before conducting any further medical-facility searches?				
Yes	60.9%	53.0%	65.0%	70.4%
No	39.1%	47.0%	35.0%	29.6%
Q6. **If No**, what was your motivation for further medical-facility searching?	***N*=59**	***N*=31**	***N*=7**	***N*=8**
Exploring procedure or treatment options at a single facility	16.9%	16.1%	28.6%	0.0%
Comparing different facilities	50.8%	54.8%	57.1%	87.5%
Comparing/finding different experts	42.4%	35.5%	42.9%	62.5%
Appointment scheduling or pre-registration	20.3%	22.6%	14.3%	12.5%
Support services (e.g., nurse lines)	3.4%	3.2%	0.0%	0.0%
Other (please specify)	10.2%[2]	9.7%	0.0%	0.0%

Notes: 1. Other explanations: Procedures offered and specialties of the medical professionals (9); Information: address, driving directions, contact information (8); Reviews (6); insurance coverage (2); scheduling (1).
2. Other explanations given: feedback on services offered (2); researching skills of medical professionals (2).

Arranging the Consultation or Treatment

Also important is how people decide that they need to visit a medical facility and if they do, what criteria affect the choice of facility. Table 2 summarizes participant responses on these decisions. We excluded the 5% of respondents who did not visit the facility for medical reasons. The "All Visits" column for Question 1 in Table 2 shows that respondents visited the facility on their own accord (47%) or as a result of advice from a medical professional (38%). Participants either visited the facility immediately after the symptoms appeared (28%) or more than two weeks had passed (38%) (Q2). The most important factors in decisions about which medical facilities to visit were proximity (53%) (described earlier in the discussion about distance to facility), insurance coverage (39%), and reputation (31%) (Q3). Interestingly, people with lower domain knowledge were more concerned about insurance coverage, perhaps reflecting a poorer understanding of how insurance networks operate and what is covered as part of insurance plans.

Table 2: Setting up consultation or treatment at medical facilities.
Statistical significance between the groups of interest is denoted by ○ $p < .05$; ● $p < .01$.

Question and Response Options	All Visits (*N*=465)	Conditional Responses		
		First Visit (*N*=114)	Domain Knowledge	
			High (*N*=59)	Low (*N*=68)
Q1. Who recommended that a consultation or treatment was required?				
No-one	46.9%	43.9%	32.2% ●	51.5%
Relative, friend, or colleague	12.7%	17.5%	10.2%	13.2%
Medical professional	37.6%	36%	55.9% ●	32.4%
Other (please specify)	2.8%[1]	2.6%	1.7%	2.9%
Q2. How soon after the symptoms emerged did it become apparent to you that consultation or treatment might be necessary?				
Immediately	27.7%	28.9%	14.6% ○	26.6%
1 day	10.1%	10.5%	16.3%	10.3%
1 week	16.1%	16.7%	20.6%	20.6%
2 weeks	8.2%	6.1%	11.2%	5.9%
Longer	37.8%	37.7%	37.3%	36.8%
Q3. What were the criteria that you used to decide which medical facility to visit? (Select one or more reasons)				
Proximity	53.1%	53.5%	50.8%	57.4%
Specialties/services offered	34.2% ○	44.7%	37.3% ○	27.9%
Particular medical professional that you found out about	23.9%	20.2%	22.0%	27.9%
Reputation	31.2%	34.2%	40.7%	39.7%
Insurance coverage	39.1%	38.6%	30.5% ●	45.6%
Primary care provider	11.8%	0.4%	12.7%	10.7%
Other (please specify)	12.9%[2]	6.1%	11.1%	5.3%

Notes: 1. Other criteria: annual checkup (*N*=2); vaccination (2); external requirement (immigration/insurance) (2); Web (1)
2. Other criteria include: recommendation from friends or family (*N*=16); referrals from doctor (7); prior experience with the facility (7), or attributes of the facility such as location (7), opening hours (in particular evenings and weekends) (3), average wait time (3), and ease of record sharing (3).

Studying the effect of domain knowledge in more detail we see that respondents with low domain knowledge were much more likely to self-diagnose (Q1: 52% Low vs. 32% High, $Z = 2.31$, $p = 0.009$), to assume that symptoms required immediate medical attention (Q2: 27% Low vs. 15% High, $Z = 1.66$, $p = 0.049$). One explanation is that people with lower domain knowledge may be more likely to leap to conclusions regarding HU, and we also see that reflected in the short time lag between HUI queries and visits (Table 1, Question 3). In contrast, those respondents reporting high medical knowledge were more likely to rely on medical professionals to recommend they seek future consultation or treatment (Q1: 56% High vs. 32% Low, $Z = 2.66$, $p = 0.003$).

To learn more about respondents' motivations, we included an open question asking: "What put you 'over the threshold' for the visit?" The responses primarily targeted symptoms (69%), especially those associated with prolonged pain or discomfort or information found by searching the Web (22%). In this group, of these respondents with low medical knowledge, 31% reported being affected in their utilization decisions by Web searches (e.g., "Web search and some research online made me believe that my symptoms were worth taking to my doctor"). In contrast, of more knowledgeable users, only 14% reported being influenced by what they found via Web search.

Search Behavior Before and After Visit to a Healthcare Professional

Recall that we are interested in the relationship between the survey and logged search behavior. One way we can do this is by asking survey respondents about health search behavior before and after the visit. Approximately 15% of survey respondents reported noticing changes in their search behavior before and after the visit. Table 3 summarizes responses regarding changes in symptom searching, serious condition searching, benign explanation searching, and searches for anything health related, for "All Visits". There were too few responses for the "First Visit" and "Domain Knowledge" conditionalizations for us to report results on those in the table.

Table 3: Percentage of users exhibiting different types of medical Web search behavior before/after the visit.

	Before Visit			After Visit		
	Increase	**Same**	**Decrease**	**Increase**	**Same**	**Decrease**
All Visits	***N*=69**			***N*=70**		
Symptoms (e.g., headache, nausea)	81.2%	17.4%	1.4%	25.7%	44.3%	30.0%
Serious conditions (e.g., brain tumor)	30.4%	62.3%	7.2%	19.9%	36.7%	43.4%
Benign explanation (e.g., headache from tension)	33.3%	65.2%	1.4%	57.6%	28.5%	13.9%
Anything health-related	20.3%	73.9%	5.8%	15.7%	74.3%	10.0%

The table shows that most respondents (81%) observed an increase in their searches for symptoms in the time before the visit to the medical facility. Other aspects of health-related searching were many times more likely to increase than decrease. Following the visit, respondents reported decreased searching for serious conditions (43% of respondents reported a drop) (as we see in Figure 1), and increased searching for benign explanations (58% of respondents reported an increase), aligning with the sharp increase in benign explanations reported in Figure 1. To further understand the reasons behind the changes, we asked survey respondents to explain the changes they reported.

Before Visit: Participants reported conducting general research (e.g., "When I got the symptoms I started searching on various other issues that I noticed ailed me, like random aches or pains"). Participants also reported that they were focused on searching on symptoms and diagnosis of potential conditions (e.g., "Searched for what could be causing symptoms from serious to less serious to normal conditions"; "Tried to figure out what the symptoms might mean. This normally sends one in a web-hopping frenzy"). There were some indications about the futility of medical searching (e.g., "…quickly discovered that web sites in general only give crazy scare tactics and tell you "Consult your doctor" so I quit reading them").

After Visit: Participants focused on searches for conditions (e.g., "I searched for info about the disease the doctor discussed with me during the appointment") or treatment options (e.g., "I looked up medications and side effects before deciding whether to fill a prescription"). Respondents also reported performing background research related to the visit (e.g., "I sought clarification on recommended course of action, wanted to find more information about my health condition, and tried to explore various options that others tried"). Respondents were asked to describe any changes in their health-related searching that they noticed following the visit. Their responses were associated with a reduction in their medical searching (e.g., "Once the doctor saw me and clarified the reasons behind my symptoms my health-related searches diminished, except for one search where I looked into the medication I was given"), focused on treatment (e.g., "I wanted to find more information about my health condition and potential home remedies"),

increased focus on diagnosed condition (e.g., "[searches] became more specific to the possible conditions associated with the diagnostic procedure I scheduled"), or an overall reduction in medical searching (e.g., "[visit] confirmed my earlier search queries … I didn't feel additional searching was necessary").

Discussion

We have presented a study seeking a better understanding of decisions about HU via the analysis of search logs and data from a user study. The survey showed that 40% of respondents performed a search about healthcare providers or facilities before engaging with them. Our contributions include a better understanding of the ways in which people decide to utilize healthcare resources and promising insights on the presence of signals in their logged search behavior that provide strong indications of HU.

The search log analysis showed a four-week gap between the emergence of searches on the symptoms and the first HUI. In addition, before the first HUI, we observe an increase in the likelihood that a concern about serious illness will be displayed (see Appendix), followed by a lift in the frequency of symptoms and disease explorations. These findings resonate with the survey responses: 40% of survey respondents waited at least two weeks after the initial concern emerged before starting the process of seeking consultation or treatment (Table 2, Question 2). The survey also showed that most people (78%) reported visiting a facility within one week of searching for it (Table 1, Question 4). Of the 15% of survey respondents who noticed changes in their medical search behavior associated with the visit, over 80% asserted an increase in the volume of symptom searching beforehand and 20-30% noticed an increase in condition searching. Survey respondents reported that during this time they searched more for symptoms, practiced self-diagnosis, and became more anxious about medical outcomes.

The logs also reveal that in the time following the input of the first HUI, there is a steep increase in searches on benign explanations after around two weeks and an associated drop in symptom searches, except for a pulse which seems to come near the steep rise in benign explanations. We hypothesize that these observations may be linked to a visit to a medical facility and participants searching for conditions highlighted to them during the visit (as the responses to the "After Visit" open questions in the survey suggest). As highlighted above, most respondents reported visiting the facility within a week or two of their HUI query. This would likely place their visit near the steep rise in benign explanation searching observed in the logs. The survey responses also showed a clear increase in benign-explanation searching and a clear decrease in serious-condition searching following the visit (Table 3). This may reflect the influence of information received during the visit. The changes in symptom searches over time can be attributed to growing or diminishing user uncertainty. In that scenario, uncertainty may return eventually as a function of time away from medical care, and is associated with another spike in HUI queries, perhaps connected once again with uncertainty resolution (see Appendix, +40 days and higher).

We cannot confirm with certainty potential conceptual links between the survey results and the observations of the behavior of populations of users online. However, the online findings may indeed be correlates of users' in-world HU behaviors, perhaps heralding receipt of online reassurance or concerning diagnosis. Patterns and trends in medical search behavior can support the development of predictive models to anticipate when users will or should seek real-world medical care or analyze search histories to build richer models for personalization or recommendation.

Limitations

We acknowledge several limitations of our study. First, survey respondents were drawn from within Microsoft, and therefore may not be fully representative of the general population. In addition, while we can make inferences about people's intentions from observing their search behavior, we cannot be certain from examining logs that they visited medical facilities. One way to address this is to study geo-coded data which contain the physical location from where log events were generated (e.g., were they proximal to a hospital). White and Horvitz (2013) explored the use of geo-coded search queries to identify visits to healthcare facilities. However that requires monitoring via GPS or cellphone triangulation that limits applications to those with access to physical location details on mobile devices. In contrast, the findings of this study show that the *effect* of HU on search behavior is observable and we may not need to track physical locations to identify potential HU episodes, only long-term search behavior (as in (White and Horvitz, 2012)), from any device, mobile or immobile.

Conclusion

We have explored the concerns, motivations, and decisions about the pursuit of professional healthcare assistance. We examined data from surveys and logs of online search activity. We believe that there is opportunity to perform richer analysis with consenting patients and to explore potential applications that might provide users with information better tailored to their needs. Avenues for future work include the use of HU estimates (e.g., the spike in benign explanations following HUI queries) as a strong signal of searcher interest in the facility sought or health care more generally. This enables search providers to incorporate evidence of *utilization*, rather than just intention inferred from HUI queries, into models of searcher interests. Estimating HU from logs also has utility when fine-grained user geolocation information is unavailable due to privacy concerns or device immobility (e.g., on desktop computers). Location may not be tracked in those cases, but from studying the *effect* that the visit has on pre- and post-search behavior, we can make inferences about when the visit occurred and its outcome—benign explanation or serious illness—based on post HU behavior. This information could be used to adapt search engine output before and after the HU, both tailored to each user and across all searchers based on aggregated search patterns.

REFERENCES

1. Asmundson, G.J.G., Taylor, S., and Cox, B.J. (2001). *Health Anxiety: Clinical and Research Perspectives on Hypochondriasis and Related Conditions*. Chichester, Wiley.
2. Ayers, S. and Kronenfeld. J. (2007). Chronic illness and health-seeking information on the Internet. *Health*, 11(3): 327–347.
3. Baker, L., Wagner, T.H., Singer, S., and Bundorf, M. (2003). Use of the internet and e-mail for health care information. *Journal of the American Medical Association*, 289(18): 2401–2406.
4. Beeferman, D. and Berger, A. (2000). Agglomerative clustering of a search engine query log. *Proc. ACM SIGKDD*, 407–416.
5. Cartright, M., White, R.W., and Horvitz, E. (2011). Intentions and attention in exploratory health search. *Proc. ACM SIGIR*, 65–74.
6. Cline, R.J. and Haynes, K.M. (2001). Consumer health information seeking on the Internet; the state of the art. *Health Education Research*, 16(6): 671–692.

7. Downey, D., Dumais, S.T., and Horvitz, E. (2007). Models of searching and browsing: languages, studies, and applications. *Proc. International Joint Conference on Artificial Intelligence,* 2740–2747.
8. Dupret, G. and Piwowarski, B. (2008). A user browsing model to predict search engine click data from past observations. *Proc. ACM SIGIR*, 331–338.
9. Eastin, M.S. and Guinsler, N.M. (2006). Worried and wired: effects of health anxiety on information-seeking and health care utilization behaviors. *Cyberpsychology and Behavior*, 9(4): 494–498.
10. Eysenbach, G. and Kohler, C. (2002). How do consumers search for and appraise health information on the World Wide Web? Qualitative studies using focus groups, usability test, and in-depth interviews. *British Medical Journal*, 324: 573–577.
11. Fox, S. and Duggan, M. (2013). Health Online 2013. *Pew Internet and American Life Project*. Available from http://pewinternet.org/Reports/2013/Health-online.aspx (Downloaded February 2013).
12. Joachims, T. (2002). Optimizing search engines using clickthrough data. *Proc. ACM KDD*, 133–142.
13. Hart, A., Henwood, F., and Wyatt S. (2004). The role of the internet in patient-practitioner relationships: findings from a qualitative research study. *Journal of Medical Internet Research*, 6(3): e36.
14. Lau, T. and Horvitz, E. (1999). Patterns of search: analyzing and modeling web query refinement. *Proc. UM*, 119–128.
15. Richardson, M. (2009). Learning about the world from long-term query logs. *ACM Transactions on the Web*, 2(4): 21.
16. Sillence, E., Briggs, P., Harris, P., and Fishwick, L. (2006). Changes in online health usage over the last 5 years. *Proc. ACM SIGCHI*, 1331–1336.
17. Tan, B., Shen, X., and Zhai, C. (2006). Mining long-term search history to improve search accuracy. *Proc. ACM SIGKDD*, 718–723.
18. White, R.W. and Horvitz, E. (2009a). Experiences with web search on medical concerns and self-diagnosis. *Proc. AMIA Annual Symposium*, 696–700.
19. White, R.W. and Horvitz, E. (2009b). Cyberchondria: studies of the escalation of medical concerns in web search. *ACM Transactions on Information Systems*, 27(4): 23.
20. White, R.W. and Horvitz, E. (2010). Web to world: Predicting transitions from self-diagnosis to the pursuit of local medical assistance in web search. *Proc. AMIA Annual Symposium*, 882–886.
21. White, R.W. and Horvitz, E. (2013). From web search to healthcare utilization: privacy-sensitive studies from mobile data. *JAMIA,* 20(1): 61-68.
22. White, R.W. and Horvitz, E. (2012). Studies of the onset and persistence of medical concerns in search logs. *Proc. ACM SIGIR*, 265-274.

Appendix

Per-day changes in symptom, condition, and HUI querying (*N* = 9,740)